

Amendments to the Claims:

This listing of claims will replace all prior versions and listings of claims in the application:

Listing of Claims:

1. (currently amended) In a computing environment, a method for implementing an association among a first data object and a second data object, the method comprising ~~the steps~~ of:

forming a first association fragment which comprises information relevant to accessing the first data object; and

forming a second association fragment which comprises information relevant to accessing the second data object, wherein the first association fragment and second association fragment cooperate to cause an association to be effectively formed between the first data object and the second data object.

2. (currently amended) The method of claim 1 wherein the first association ~~part~~ fragment and the second association part ~~association fragment~~ are maintained in the computing environment as data objects apart from, and in like manner as peers to, the first data object and the second data object.

3. (currently amended) The method of claim 1 wherein the first data object and the second data object are maintained by separate processes in the computing environment.

4. (currently amended) The method of claim 1 wherein the first data object and the second data object are maintained in separate data stores.

5. (currently amended) The method of claim 1 wherein the first association fragment is maintained substantially in proximity to the first data object to facilitate ~~[[the]]~~ an interaction therebetween, and wherein the second association fragment is maintained substantially in proximity to the second data object to facilitate ~~[[the]]~~ an interaction therebetween.

6. (currently amended) In a computing environment, a method for implementing an association between a first data object and a second data object, the method comprising ~~the steps~~ of:

forming a first association fragment which comprises information relevant to accessing the first data object; and

forming a second association fragment which comprises information relevant to accessing the second data object, wherein:

the first association fragment and second association fragment cooperate to cause an association to effectively be formed between the first data object and the second data object,

the first data object and the first association fragment are located in a first partition of storage, and

the second data object and the second association fragment are located in a second partition of the storage.

7. (currently amended) A data processing system implemented method managing associations among a plurality of data objects, said plurality of data objects are categorized by data type and associated with data stored in a plurality of location subdivisions based on the data type, the method comprising:

receiving a first request for a data object associated first data object from a client, said request includes a first data object identifier, wherein said first data object is categorized as being a first data type;

identifying a first location subdivision for the first data object identifier;

accessing first fragmented association information corresponding with said first location subdivision, wherein the first fragmented association information includes association information related to data objects stored in said first location subdivision;

finding, in the first fragmented association information, a data object associated with said first data object based on the first data object identifier; and

returning information related to the data object or objects associated with said first data object to the client.

8. (currently amended) The method recited in claim 7, wherein the data type comprises entity information.

9. (currently amended) The method recited in claim 7 further comprises:

receiving a second request for a data object associated second data object from a client, said second request including a second data object identifier, wherein said first data object and second data ~~objects~~ object are categorized as being the first type data;

identifying a second location subdivision for the second data object identifier;

accessing second fragmented association information corresponding with said second location subdivision, wherein the second fragmented association information includes association information related to data objects stored in said second location subdivision;

finding, in the second fragmented association information, a data object associated with said second data object based on the second data object identifier; and

returning information related to the data object associated with said second data object to the client.

10. (currently amended) The method recited in claim 7, wherein the information related to the data object associated with said first data object is an instance of the data object associated with said first data object.

11. (previously presented) The method recited in claim 9, wherein the first fragmented association information corresponding with said first location subdivision, and the second fragmented association information corresponding with said second location subdivision

information are stored external to respective said first location subdivision and said second location subdivision.

12. (currently amended) The method recited in claim 11, wherein the first fragmented association information and the first location subdivision reside in a first proximity domain.

13. (previously presented) The method recited in claim 12, wherein the second fragmented association information and the second location subdivision information reside in a second proximity domain.

14. (previously presented) The method recited in claim 13, wherein said first proximity domain and said second proximity domain are overlapping.

15. (previously presented) The method recited in claim 9, wherein a first partition service manages data objects for the first location subdivision and a second partition service manages data objects for the second location subdivision.

16. (currently amended) The method recited in claim 7, wherein said request for a data object associated first data object from a client further includes a requested data type, and

finding in the first fragmented association information a data object associated with said first data object based on the first data object identifier further comprises:

traversing the first fragmented association information based on the requested data type and the first data object identifier.

17. (currently amended) The method recited in claim 7, wherein an association service receives the request for [[a]] the data object associated first data object from [[a]] the client.

18. (currently amended) The method recited in claim 17, wherein the first data object identifier is a unique primary key (PK) and identifying a first location subdivision for the first data object identifier further comprises:

accessing a plurality of primary key (PK)-location subdivision maps held in a satellite cache, wherein said satellite cache and said association service are in a local domain; and

looking up a location subdivision identifier for the first location subdivision in the plurality of PK-location subdivision maps from the satellite cache based on the unique PK for the first data object.

19. (currently amended) The method recited in claim 18, wherein:

said looking up the location subdivision identifier in the satellite cache using the plurality of PK-location subdivision maps returns a null set of location subdivision identifiers[[,]] and said identifying a first location subdivision for the first data object identifier further comprises:

accessing a plurality of PK-location subdivision maps held in an enterprise repository, wherein said enterprise repository is external to the local domain; and

looking up the location subdivision identifier for the location subdivision holding the data object in the plurality of PK-location subdivision maps in the enterprise repository based on the unique PK for the first data object.

20. (currently amended) The method recited in claim 19, wherein said looking up the location subdivision identifier in the enterprise repository using the plurality of PK-location subdivision maps returns a null set of location subdivision identifiers identifying a location subdivision holding the first data object, and the method further comprises:

finding a steward;

requesting [[the]] an identity of the location subdivision holding the first data object from the steward, wherein said steward issues a query to each location subdivision for the identity of the location subdivision holding the first data object; and

returning the identity of the location subdivision holding the first data object to the association service.

21. (previously presented) The method recited in claim 20, wherein the steward manages data objects of the first data type.

22. (currently amended) The method recited in claim 7, wherein the information related to the data object associated with said first data object comprises instances of all data objects associated with the first data object identifier.

23. (currently amended) A data processing system implemented method for managing associations among a plurality of data objects, the method comprising:

categorizing [[a]] the plurality of data objects by data type;

subdividing ones of the data objects of a first data type into a first plurality of subdivisions;

subdividing ones of the data objects of a second data type into a second plurality of subdivisions;

storing the ones of the data objects in the first plurality of subdivisions in a respective first plurality of location subdivisions;

storing the ones of the data objects in the second plurality of subdivisions in a respective second plurality of location subdivisions;

defining associations between each of the data ~~object~~ objects in each of the first plurality of location subdivisions and all of the data objects in each of the second plurality of location subdivisions;

forming a fragmented association record corresponding with each location subdivision of the first plurality of location subdivisions, wherein each of the fragmented association ~~record~~ records comprises association information related to data objects stored in ~~[[a]]~~ one of the first plurality of location subdivisions;

defining associations between each of the data objects in each of the second plurality of location subdivisions and all of the data objects in each of the first plurality of location subdivisions; and

forming a fragmented association record corresponding with each location subdivision of the second plurality of location subdivisions, wherein each of the fragmented association ~~record~~ records comprises association information related to data objects stored in ~~[[a]]~~ one of the second plurality of location subdivisions.

24. (currently amended) The method recited in claim 23, wherein each of said fragmented association ~~record~~ records is stored external to the corresponding location subdivision.

25. (currently amended) The method recited in claim 23, wherein ~~[[a]]~~ one of the fragmented association ~~record~~ records corresponding with one of the first plurality of location

subdivisions ~~contains~~ includes association information defining a plurality of associations for one of the data objects.

26. (currently amended) The method recited in claim 23, wherein the data type comprises entity information.

27. (currently amended) A data processing system implemented program product embodied on a processing system readable medium for implementing a method for managing associations among a plurality of data objects, said plurality of data objects are categorized by data type and associated with data stored in a plurality of location subdivisions based on the data type, the program product comprising:

instructions for receiving a request for a data object associated first data object from a client, said request includes a first data object identifier, wherein said first data object is categorized as being a first data type;

instructions for identifying a first location subdivision for the first data object identifier;

instructions for accessing first fragmented association information corresponding with said first location subdivision, wherein the first fragmented association information includes association information related to data objects stored in said first location subdivision;

instructions for finding, in the first fragmented association information, a data object associated with said first data object based on the first data object identifier; and

instructions for returning information related to the data object or objects associated with said first data object to the client.

28. (currently amended) The program product recited in claim 27, wherein the data type comprises entity information.

29. (currently amended) The program product recited in claim 27 further comprises:

instructions for receiving a second request for a data object associated second data object from ~~[[a]]~~ the client, said second request includes a second data object identifier, wherein said first data object and said second data object ~~objects~~ are categorized as being the first data type;

instructions for identifying a second location subdivision for the second data object identifier;

instructions for accessing second fragmented association information corresponding with said second location subdivision, wherein the second fragmented association information includes association information related to data objects stored in said second location subdivision;

instructions for finding in the second fragmented association information a data object associated with said second data object based on the second data object identifier; and

instructions for returning information related to the data object associated with said second data object to the client.

30. (currently amended) The program product recited in claim 27, wherein the information related to the data object associated with said first data object is an instance of the data object associated with said first data object.

31. (previously presented) The program product recited in claim 29, wherein the first fragmented association information corresponding with said first location subdivision, and the second fragmented association information corresponding with said second location subdivision information are stored external to respective said first location subdivision and said second location subdivision.

32. (previously presented) The program product recited in claim 31, wherein the first fragmented association information and the first location subdivision reside in a first proximity domain.

33. (previously presented) The program product recited in claim 32, wherein the second fragmented association information and the second location subdivision information reside in a second proximity domain.

34. (previously presented) The program product recited in claim 33, wherein said first proximity domain and said second proximity domain are first and second multicast domains.

35. (previously presented) The program product recited in claim 29 further comprises:

instructions for implementing a first partition service for managing data objects for the first location subdivision; and

instructions for a second partition service for managing data objects for the second location subdivision.

36. (currently amended) The program product recited in claim 27, wherein:

said request for a data object associated first data object from a client further includes a requested data type, and

the instructions for finding, in the first fragmented association information, a data object associated with said first data object based on the first data object identifier further comprises:

instructions for traversing the first fragmented association information based on the requested data type and the first data object identifier.

37. (currently amended) The program product recited in claim 27 further comprises;

instructions for implementing an association service for executing the instructions for receiving the request for a data object associated first data object from a client.

38. (currently amended) The program product recited in claim 37, wherein the first data object identifier is a unique primary key (PK) and the instructions for identifying a first location subdivision for the first data object identifier further comprises:

instructions for implementing a satellite cache service for holding a plurality of primary key (PK)-location subdivision maps, wherein said satellite cache service and said association service are implemented in a local domain;

instructions for accessing the plurality of primary key (PK)-location subdivision maps held by the satellite cache service; and

instructions for looking up a location subdivision identifier for the first location subdivision in the plurality of PK-location subdivision maps from the satellite cache service based on the unique PK for the first data object.

39. (currently amended) The program product recited in claim 38 wherein the instructions for identifying a first location subdivision for the first data object identifier further comprises:

instructions for implementing an enterprise repository service for holding and accessing a plurality of PK-location subdivision maps, wherein executing instructions for implementing said enterprise repository is performed externally to the local domain; and

instructions for looking up the location subdivision identifier for the location subdivision holding the data object in the plurality of PK-location subdivision maps in the enterprise repository service based on the unique PK for the first data object in response to executing the instructions for looking up the location subdivision identifier in the satellite cache service using the plurality of PK-location subdivision maps ~~returns~~ and returning a null set of location subdivision identifiers.

40. (currently amended) The program product recited in claim 39, wherein the instructions for looking up the location subdivision identifier in the enterprise repository using the plurality of PK-location subdivision maps further comprises:

instructions for finding a steward service in response to executing the instructions for looking up the location subdivision identifier in the enterprise repository ~~returns~~ and returning a null set of location subdivision identifiers identifying a location subdivision holding the first data object;

instructions for requesting the identity of the location subdivision holding the first data object from the steward service, wherein said steward service issues a query to each location subdivision for the identity of the location subdivision holding the first data object; and

instructions for returning the identity of the location subdivision holding the first data object to the association service.

41. (previously presented) The program product recited in claim 40, wherein the steward service manages data objects of the first data type.

42. (currently amended) The program product recited in claim 27, wherein the information related to the data objects associated with said first data object comprises instances of all of the data objects associated with the first data object identifier.

43. (currently amended) A data processing system implemented program product embodied on a data processing system ~~processing system~~ readable medium for implementing a method for managing associations among a plurality of data objects, the program product comprising:

instructions for categorizing a plurality of data objects by data type;

instructions for subdividing the data objects of a first data type into a first plurality of subdivisions;

instructions for subdividing the data objects of a second data type into a second plurality of subdivisions;

instructions for storing the data objects in the first plurality of subdivisions in a respective first plurality of location subdivisions;

instructions for storing the data objects in the second plurality of subdivisions in a respective second plurality of location subdivisions;

instructions for defining associations between each of the data ~~object~~ objects in each of the first plurality of location subdivisions and all of the data objects in each of the second plurality of location subdivisions;

instructions for forming a fragmented association record corresponding with each of the location ~~subdivision~~ subdivisions of the first plurality of location subdivisions, wherein each of the fragmented association records ~~record~~ comprises association information related to ones of the data objects stored in [[a]] one of the first plurality of location subdivisions;

instructions for defining associations between each of the data objects in each of the second plurality of location subdivisions and all of the data objects in each of the first plurality of location subdivisions; and

instructions for forming a fragmented association record corresponding with each of the location ~~subdivision~~ subdivisions of the second plurality of location subdivisions, wherein each of the fragmented association records ~~record~~ comprises association information related to ones of the data objects stored in [[a]] one of the second plurality of location subdivisions.

44. (currently amended) The program product recited in claim 43, wherein each of said fragmented association ~~record~~ records is stored external to the corresponding location subdivision.

45. (currently amended) The program product recited in claim 43, wherein a fragmented association record corresponding with one of the first plurality of location subdivisions ~~contains~~ includes association information defining a plurality of associations for one of the data objects.

46. (currently amended) The program product recited in claim 43, wherein the data type comprises entity information.

47. (currently amended) A system for managing associations among a plurality of data objects, the system comprising~~[[:]]~~ a plurality of network domains, wherein each of the network domains comprises:

a first data processing system of ~~[[the]]~~ a plurality of multipurpose data processing systems including a first processor and a first memory for running a first partition service, said first partition service being configured to manage ~~manages~~ data of a first data type stored locally in said memory, ~~wherein said data is a first type~~;

a second data processing system of the plurality of multipurpose data processing systems including a processor and a memory for running a second partition service, said second partition

service being configured to manage ~~manages~~ data of a second data type stored locally in said second memory; ~~wherein said data is a second type~~; [[and]]

a third data processing system including a processor and a memory for running a satellite service, said satellite service being configured to maintain ~~maintains~~ a plurality of primary key (PK)-partition maps for identifying a partition service managing a data object by a unique PK;

a fourth data processing system including a processor and a memory for running a first fragmented association engine information corresponding with said first partition service, said first fragmented association engine being configured to manage ~~manages~~ association information related to data objects managed by said first partition service; [[and]]

a fifth data processing system including a processor and a memory for running a second fragmented association engine information corresponding with said second partition service, said second fragmented association engine being configured to manage ~~manages~~ association information related to data objects managed by said second partition service; and

a sixth data processing system including a processor and a memory for running a steward service, said steward service being configured to manage ~~manages~~ a plurality of PKs for [[a]] the first data type.

48. (currently amended) The system recited in claim 47, wherein said third data processing system [[is]] includes one of said first data processing system [[and]] or said second data processing system.

49. (currently amended) The system recited in claim 47, wherein said steward service is configured to issue ~~issues~~ a query to said first partition service and said second partition service for an identity of a partition service managing a data object based on a unique PK.

50. (currently amended) The system recited in claim 47, wherein the first data type and the second data type ~~comprises~~ comprise entity information.

51. (currently amended) The method recited in claim 47, wherein each of the network domains comprises a data processing system for running an association service, said association service ~~comprises~~ being configured to:

~~receives~~ receive a request for a data object from a client, said request includes a data object identifier for an identified data object;

~~identifies~~ identify a third partition service based on the data object identifier for the identified data object;

~~references~~ reference a fragmented association engine corresponding with said third partition service, said fragmented association engine manages association information related to data objects managed by said third partition service;

~~finds~~ find a fourth partition service based on said fragmented association engine corresponding with said third partition service; and

~~returns~~ return information related to the requested data object associated with said identified data object to the client.

52. (currently amended) The system recited in claim 51, wherein the information related to the requested data object associated with said identified data object is an instance of the requested data object.

53. (previously presented) The system recited in claim 47, wherein the first fragmented association engine corresponding with said first partition service, and the second fragmented association engine corresponding with said second partition service are stored external to respective said first partition service and said second partition service.

54. (previously presented) The system recited in claim 47, wherein the first fragmented association service and the first partition service reside in a first proximity network domain.

55. (previously presented) The system recited in claim 47, wherein the second fragmented association engine and the second partition service reside in a second proximity network domain.

56. (previously presented) The system recited in claim 55, wherein said first proximity network domain and said second proximity network domain are overlapping.

57. (currently amended) The system recited in claim 51, wherein:
said request for a data object associated with said data object from a client further includes a requested data type, and
said fragmented association engine traverses said association information related to data objects managed by said third partition service based on the requested data type and the data object identifier.

58. (currently amended) The system recited in claim 51, wherein the data object identifier ~~[[is]]~~ includes a unique primary key (PK) and said association service is configured to identify ~~identifies a~~ the third partition service based on the data object identifier for the identified data object using said satellite service, wherein said satellite service and said association service are in a local network domain, and said satellite service is configured to look ~~looks up a~~ partition service identifier for the third partition service in the plurality of PK-partition maps from the satellite service based on the unique PK for the identified data object.

59. (currently amended) The system recited in claim 58 further comprises:

a seventh data processing system including a processor and a memory for running an enterprise service, said enterprise service being configured to manage ~~manages~~ a plurality of PK-partition service maps for the plurality of network domains, wherein a partition service identifier for a partition service managing said identified data object is identified in the plurality of PK-partition service maps for the plurality of network domains based on the unique PK for the identified first data object in response to said satellite service returning a null set of partition service identifiers.

60. (currently amended) The system recited in claim 59, wherein in response to the enterprise repository returning [[a]] the null set of partition identifiers, said association service utilizes said steward service for identifying a partition service managing the identified data object.

61. (currently amended) The system recited in claim 60, wherein the steward is configured to manage ~~manages~~ data objects of the first data type.

62. (currently amended) The system recited in claim 51, wherein the information related to the requested data object associated with said identified data object comprises instances of all data objects associated with the data object identifier.

63. (new) The method of claim 1, wherein:

the first association fragment and the first data object are co-located in a first subdivision of computer storage, and

the second association fragment and the second data object are co-located in a second subdivision of computer storage.

64. (new) The method of claim 63, further comprising:

traversing a relationship from one of the first data object or the second data object by using one of the association fragments coupled to the one of the first data object or the second data object.